

10-02-00

K4

PATENT APPLICATION TRANSMITTAL LETTER

(Large Entity)

Docket No.

NCRC-0028-US (9433)

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Transmitted herewith for filing under 35 U.S.C. 111 and 37 C.F.R. 1.53 is the patent application of:

ALEX P. YUNG, GREG KORENEVSKY, AND STEVEN FEINHOLZ

For: METHOD AND APPARATUS FOR PERFORMING PARALLEL DATA OPERATIONS

Enclosed are:

- ☒ Certificate of Mailing with Express Mail Mailing Label No. EL594060551US
- ☒ ELEVEN sheets of drawings.
- ☐ A certified copy of a application.
- ☐ Declaration ☐ Signed. ☐ Unsigned.
- ☐ Power of Attorney
- ☐ Information Disclosure Statement
- ☐ Preliminary Amendment
- ☐ Other:

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	42	- 20 =	22	x \$18.00	\$396.00
Indep. Claims	5	- 3 =	2	x \$78.00	\$156.00
Multiple Dependent Claims (check if applicable) <input type="checkbox"/>					\$0.00
BASIC FEE					\$690.00
TOTAL FILING FEE					\$1,242.00

- ☐ A check in the amount of to cover the filing fee is enclosed.
- ☐ The Commissioner is hereby authorized to charge and credit Deposit Account No. as described below. A duplicate copy of this sheet is enclosed.
- ☐ Charge the amount of as filing fee.
- ☐ Credit any overpayment.
- ☐ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

Dated: SEPTEMBER 29, 2000



Signature

DAN C. HU
 REGISTRATION NO. 40,025
 TROP, PRUNER & HU, P.C.
 8554 KATY FRWY, SUITE 100
 HOUSTON, TEXAS 77024
 PHONE: (713) 468-8880
 FAX: (713) 468-8883
 CUSTOMER NO. 21906

CC:

APPLICATION

UNITED STATES LETTERS PATENT

005250 2552960

TITLE: METHOD AND APPARATUS FOR PERFORMING PARALLEL
DATA OPERATIONS

INVENTORS: ALEX P. YUNG, GREG KORENEVSKY, AND STEVEN
FEINHOLZ

Express Mail No.: EL594060551US

Date: September 29, 2000

Prepared by: Trop, Pruner & Hu, P.C.
8554 Katy Freeway, Ste. 100, Houston, TX 77024
713/468-8880 [Office], 713/468-8883 [Fax]

METHOD AND APPARATUS FOR PERFORMING PARALLEL DATA OPERATIONS

TECHNICAL FIELD

The invention generally relates to storing, filtering, and accessing data in a parallel manner.

BACKGROUND

5 One of the key advantages of storing large amounts of data in a database is that a specific subset of the stored data can be retrieved in an organized manner. Often, the subset of the stored data that is retrieved is analyzed to study various indications, such as economic trends, consumer reactions, and the like. To learn about customers, businesses are collecting various types of
10 information about their customers, such as personal data, geographic/demographic data, purchasing habits, and so forth. Such customer data are stored in a database system, such as in a relational database management system (RDBMS), where the data can be processed and sorted into a format suitable for reporting or analysis. An example of a database system in
15 which such information is collected is a data warehouse in which data is input from a variety of sources and organized into a format that is structured for query and analysis or reporting. The volume of data collected in most large data warehouses is at least several gigabytes and often exceeds tens or even
20 hundreds of terabytes.

To handle the massive amount of data that is collected and processed in such data warehouses, sophisticated platforms are typically employed. The platforms include parallel processing systems, such as massive parallel processing (MPP) systems or symmetric multiprocessing (SMP) systems. An MPP
25 system typically is a multi-node system having a plurality of physical nodes interconnected by a network. An SMP system typically is a single-node system

having multiple processors. Collected data is stored in storage devices in such systems, which are accessible by the various nodes or processors to perform processing. In a parallel system, stored data portions are accessible in parallel to increase access speeds.

5

Many times a user interfaces with a database system to implement several tasks. These tasks include storing data, retrieving data, performing data queries, and the like. In order to utilize computing resources efficiently, these tasks can be performed in parallel. The control of the tasks described above can be performed in the database, a database server system, or from a remote system such as a client system.

10

In conventional parallel processing environments, a typical application usually includes several tasks. Each of these tasks is generally responsible for a portion of an application's workload. Sometimes, an application can be parallelized based upon its functions; that is, each task can perform a different function. This process is called functional parallelism. Another way of parallelizing an application is to divide its input, output or intermediate data into multiple portions and to assign a task for each data portion. This method is often called data parallelism. In either case, each such task is usually independent of other tasks in the application in the sense that they do not need to share internal processing states or information with each other.

15

20

Consequently, each task can be executed independently of other tasks, in a concurrent or simultaneous manner, the latter being the case in systems with multiple processors

25

One of the difficulties frequently encountered in implementing a parallel application is the need to coordinate the processing of the individual parallel tasks. Currently, the most common approaches used to address this issue are broadcasting coordination requests from every task to all the other tasks and creating a central component that dictates and controls the processing and

30

communication between the tasks. While the first approach is suitable for a certain class of parallel applications, it frequently leads to increased complexity in the design of the parallel application and also introduces increased communication overhead, which can impede the application's scalability.

5

The second approach is frequently used in transactional and database systems (e.g. the 2-phase commit protocol). It, too, restricts the flexibility in the design of the individual parallel tasks in the sense that the processing in each task is dictated by a statically defined protocol (such as a fixed number of steps or phases) implemented in a controller. This static protocol, which does not change from one application to another, lacks application-specific semantics that are usually required by complex applications such as ETL (**E**xtract, **T**ransform, **L**oading) applications used in a data warehousing environment.

10

15

Conventional parallel execution of data tasks employ synchronization functions; such as WAIT, POST, LOCK, UNLOCK, GROUP, BARRIER, and the like; which are generally platform-dependent. These platform-dependent synchronization functions require multiple implementations in order for applications to run in a heterogeneous environment. Multiple implementations of the synchronization functions are less efficient and require valuable computer resources.

20

SUMMARY

25

Described below is a technique for performing parallel data operations upon data in a database. This technique includes receiving a data transaction request in a client system and executing a plurality of multi-phase parallel tasks in response to the request to perform the data operations upon the data in the database.

30

In general, in accordance with one embodiment, an apparatus for performing parallel data operations upon data in a database, comprises: a user

interface; a processor coupled with the user interface, wherein the processor receives a data transaction request from the user interface; and a controller coupled with the processor, wherein the controller performs a number of tasks in parallel based upon instructions received from the processor, each task
5 performed in a plurality of phases.

Other features and embodiments will become apparent from the following description, from the drawings, and from the claims.

10

BRIEF DESCRIPTION OF THE DRAWINGS

The invention may be understood by reference to the following description taken in conjunction with the accompanying drawings, in which like reference numerals identify like elements, and in which:

15

Figure 1 is a block diagram of an example database communications system, including a database system and a client system, in accordance with an embodiment of the present invention;

20

Figure 2 is a more detailed depiction of one embodiment of the database system illustrated in Figure 1;

25

Figure 3 is a block diagram of one embodiment of the database that can be used in the communication system of Figures 1 and 2;

30

Figure 4 is a block diagram of one embodiment of the client system illustrated in Figure 1;

Figure 5 is a block diagram of one embodiment of a client controller in the client system of Figure 4;

Figures 6A and 6B illustrate one example of a master-slave implementation of the multi-phase protocol in accordance with one embodiment of the present invention;

5 Figure 7 is a flowchart of one embodiment of the method of performing parallel data operations;

10 Figure 8 is a flowchart of one embodiment of the method of preparing to perform a multi-phase coordination protocol;

 Figure 9 is a flowchart of the method of performing the parallel data operations described in Figure 7; and

15 Figure 10 is a block diagram of another database communications system, in accordance with another embodiment of the present invention.

DETAILED DESCRIPTION

20 In the following description, numerous details are set forth to provide an understanding of the present invention. However, it will be understood by those skilled in the art that the present invention may be practiced without these details and that numerous variations or modifications from the described embodiments may be possible.

25 Generally, when data is transmitted from one device to another, such as from a server system to a client system, a certain amount of efficiency in performing data operations is desired. Embodiments of systems used to perform efficient parallel processing of data tasks between a client system and a database system are discussed below. The embodiments shown below illustrate a system that is capable of performing multi-phase synchronization of data tasks executed
30 in parallel. The embodiments shown below illustrate a plurality of statically different synchronization points that are dictated by the central entity. Parallel

implementation of data tasks are synchronized by the central entity in a multi-phase processing environment.

Figure 1 shows a system that comprises a database system 110 and a client system 120, in accordance with one arrangement of the present invention. The client system and the database system are connected together by a network connection 145. The client system 120 can be used to request a number of data-related operations in the database system 110. In the illustrated embodiment, the database system 110 receives data transaction requests from the client system 120. A data transaction request can be a request to store data, a request to retrieve data, a request to query data, a request to filter data, a request to erase data, and the like. The data transaction requests, which in one example are or Structured Query Language (SQL) queries, are executed by the database system 110. To communicate the SQL queries from the client system 120 to the database system 110, the SQL queries may be carried in requests according to a predetermined format. In one embodiment, the client system 120 performs a plurality of data tasks upon the database system 110 in parallel, via multiple paths, or sessions, to the database system 110. The network connection 145 supports a plurality of paths or sessions from the client system 120 to the database system 110.

In one embodiment, network connection 145 is typically either a private network connection or a public network connection, or a combination of both. Examples of the network connection 145 include communication networks, channels, links, or paths, and systems or devices (such as routers) used to route data over such networks, channels, links, or paths. A private network connection 145 typically includes local area networks (LANs) or wide area networks (WANs), while one example of a public network connection is the Internet. The types of devices or systems that have access to the communication networks include both wired and wireless devices or systems. Examples of wired devices include personal desktop computer systems that have access to the network connection

145 through dial-up services or other types of data network connection services. Wireless devices (e.g., portable computers with wireless modems, personal digital assistants, mobile telephones) have access to the network connection 145 over wireless infrastructures.

5

Turning now to Figure 2, a more detailed illustration of one embodiment of the database system 110, is shown. In one embodiment, the database interface 210 is capable of collecting requested or queried data from the database 230 in response to requests or queries from the client system 120. The database interface 210 is also capable of performing other data tasks, such as storing data into the database 230, retrieving data from the database 230, and the like. The database interface 210, along with the server 220, are capable of supporting multiple data operation tasks in a parallel manner. The system described in Figure 1 is capable of facilitating data operations in a parallel manner, between the database system 110 and the client system 120.

In the illustrated embodiment, the server 220 receives data transaction requests from the client system 120. The requests are translated to a format that can be understood by the database controller 210, which in turn executes the tasks directed by the client system 120. In another arrangement, the client system 120 is coupled directly to the database interface 210 without going through the server 220. The client system 120 is capable of collecting data from a plurality of data sources 250 and storing them into the database system 110. The data source 250 can be a point-of-sale terminal in a retail environment, a hand-held terminal, a centrally-located computer system, a mainframe computer system, or the like. For example, the client system 120 gathers sales transaction data from a point-of-sale terminal in a retail environment and stores the data into the database system 110.

Turning now to Figure 3, one embodiment of a database unit 300 is illustrated. The illustrated database unit 300 in one example is made up of the

database 230 and the database interface 210 (Figure 2). In the illustrated embodiment, the database unit 300 is a multi-node parallel system that has a plurality of nodes 320, 330, 340 that are coupled by an interconnect network 310. The database unit 300 also comprises a parsing engine 370, which is capable of parsing data from the plurality of nodes 320, 330, 340. In another alternative embodiment, the database unit 300 comprises the database interface 210, the database 230, and the server 220. Each node 320, 330, 340 comprises one or more storage devices 350 and at least one memory storage controller 360. The arrangement of Fig. 3 is an example of a parallel database system. In further embodiments, other types of arrangements can be employed, such as single-node systems.

Turning now to Figure 4, a more detailed illustration of one embodiment of the client system 120 is shown. The client system 120 comprises a client interface 410, a client controller 420, a client processor 430, and a user interface 440. The client interface 410 is coupled to the network connection 145 such that communications from the client system 120 and the database system 110 is enabled.

The client processor 430 receives data transaction requests and other instructions from the user interface 440. The client processor 430 then prompts the client controller 420 to execute the data transaction requests received from the user interface 440. Subsequently, the client controller 420 implements a multi-phase coordination protocol, in accordance with an embodiment, using the data transaction request.

Figure 5 illustrates a more detailed description of one embodiment of the client controller 420. The client controller 420 comprises: a task coordinator 520; a first component 530A; a second component 530B; an Nth component 530N; and a communications interface 590. In one embodiment, the components 530 are software operators or processes that can perform data

operations, such as data extract, data transform, data load, and other functions.

The task coordinator 520 implements the request received by the client processor 430. The task coordinator 520 initiates the multi-phase coordination protocol, in accordance with the present invention, in order to perform parallel tasks called for by the client processor 430. In one embodiment, the components 530 use the methods 512 and functions 512 in order to perform the parallel tasks called for by the client processor 430. In one embodiment, the methods 512 and functions 514 are UNIX statements, SQL statements, or other types of methods or functions.

When the client processor 430 receives one or more data transaction requests from the user interface 440, the data transaction requests are analyzed by the client processor 430. In one embodiment, the client processor 430 then generates a task plan 595 in order to execute the request. In an alternative arrangement, the task plan 595 is generated by an external source (not shown). The task coordinator 520 selects one or more operators or components 530, which are used to perform the operation specified by the data transaction request. The operators or components 530 implement, or invoke, their tasks using the methods 512 or functions 514. Examples of the methods 512 used by the task coordinator 520 include a "Start-up" method, a "Terminate" method, an "Execute" method, a "Checkpoint" method, a "Restart" method, and the like.

The components 530 are able to communicate with each other over the communications interface 590. Data can be input into, or output from, the components 530 via the communications interface 590. When the task coordinator 520 receives a task plan 595, the task coordinator 520 defines a task that is executed by the components 530.

The tasks defined by the task coordinator 520 are then executed as steps that run in parallel with each other. In other words, a plurality of components 530 perform the tasks defined by the task coordinator 520 in a parallel manner.

The task coordinator 520 implements the tasks using the components 530, the methods 512, and the functions 514.

The task coordinator 520 generates a multi-phase arrangement, wherein the tasks defined by the task coordinator 520 are performed within predetermined phases of the multi-phase arrangement. In one embodiment, a phase is a portion of a sequence in which one or more events called for by a task is performed. The task coordinator 520 synchronizes the tasks performed by the components 530. In one arrangement, the task coordinator 520 assigns the multi-phase structure such that each component 530 performs a task with a predetermined number or phase(s). For example, the first component 530A may perform a particular task in the first phase. Simultaneously, the second component 530B may perform another task independent of the first task. Therefore, the first and the second tasks are performed simultaneously and independently. In other words, the first and the second tasks are performed in a parallel manner. The second component 530B may require more than one phase to complete the task assigned to the second component 530B. Therefore, the task coordinator 520 will schedule a second phase in order to allow the second component 530B to complete its assigned task. The functions of the second component 530B can be performed independently and transparently to the functions of the first component 530A. Additionally, the first component 530A and the second component 530B may exchange data via the communications interface 590. The task coordinator 520 is capable of orchestrating multiple-tasks by invoking the components 530 in a plurality of phases.

In one embodiment, the first through Nth components 530 each are different instances of the methods 512. The methods 512 that can be invoked to create the components 530 include the "Startup" method, the "Execute" method, the "Checkpoint" method, the "Restart" method, and the "Terminate" method.

Referring now to Figures 6A and 6B, examples of generating the

components that operate in a multi-phase parallel mode are illustrated. The example illustrates a Fastload operation, in which data is collected from one or more data sources and loaded into the database system 110. In this example, the Startup method (corresponding to the startup phase of the Fastload operation) is invoked three times to create three components 530 (or instances) that operate in parallel. In one embodiment in which a master-slave arrangement is used, the first component 530A is designated as the master component. The second and third components 530B, 530N can be designated as the slave components. The first component 530A example performs the Startup method in three phases. The second and third components 530B, 530N, which are the slave components, need only two phases to complete their respective tasks. The master component 530A and the slave components 530B, 530N operate in parallel. During the first phase, the master component 530A and the slave components 530B, 530N create connections for SQL sessions.

At the end of the first phase, the master and slave components 530 send a "NextPhase" message to the task coordinator 520. The "NextPhase" message indicates to the task coordinator 520 that the component 530 should be invoked in the next phase. The task coordinator 520 then invokes the components 530 that returned the "NextPhase" message for another phase. During the second phase, the master component 530A and the slave components 530B, 530N create a Fastload session. At the end of the second session, the slave components 530B, 530N send an "EndMethod" message to the task coordinator 520, and the master component 530A sends a "NextPhase" message to the task coordinator 520. The "EndMethod" indicates to the task coordinator 520 that the method should not be used to invoke the component 530 that sent the "EndMethod" message. The task coordinator 520 then only invokes the master coordinator 530A for a third phase, wherein the "Begin Loading" command is issued to start the loading process, which includes retrieving data from data sources and loading the data into the database system 110. At the end of the third phase, the master component 530A issues an "EndMethod" to terminate the

operation of the master component 530. This indicates to the task coordinator 520 that none of the components 530 need to be re-invoked.

An "Execute" method is used to load data from the client system 120 into the database system 110. Like the "Startup" method, the "Execute" method is carried out by three components, a master component 530A and two slave components 530B, 530C, operating in parallel. The master component 530A performs its assigned task in two phases, while the slave components 530B, 530N need only one phase to complete their respective tasks. In the first phase, the master component 530A and slave components 530B, 530N each read a record from a data source that is to be loaded into the database system 110, until it reaches an "end-of-data" indicator. Upon reaching the end-of-data point, each component 530 issues a NextPhase message to the task coordinator 520, which then invokes the respective component 530 for a second phase. If no end-of-data indicator is detected after reading a record, the component sends a "SamePhase" message to the task coordinator 520, which causes the task coordinator 520 to invoke the same components 530 within the same phase. In this example, only the master coordinator 530A is invoked in the second phase. In one embodiment, the master coordinator 530A issues an "End Loading" command to end the loading process and sends a "EndMethod" message to the task coordinator 520, which terminates the master component 530A.

To enable the components 530 to be executed in a checkpoint-restartable fashion (*i.e.*, being able to re-enter a task-position and continue that task from the point it was previously halted), the "Checkpoint" method can be used to invoke a master component 530A and two slave components 530. The master and the slave components 530 operate in parallel. The components 530 write statistics to global variables, which provide indications of the current execution point of each component 530. If a fault were to occur for any reason, the global variables could be accessed to enable a restart from the last checkpoint. After the first phase, the master component 530A issues a "NextPhase" message to

the task coordinator 520, while the slave components 530 issue "EndMethods" to the task coordinator 520. The task coordinator 520 invokes the master component 530A for another phase, while the task coordinator 520 terminates the slave components. The master component 530A writes checkpoint records to the log and issues an "EndMethod" message to the task coordinator 520, which then terminates the master component 530A. Similarly, other methods, such as the "Restart" and the "Terminate" methods, can be used to invoke components 530 for multi-phase parallel operation.

The previous examples illustrate the synchronization achieved by the implementation of the multi-phase operation in accordance with the present invention. In some cases the slave components 530B, 530N complete their respective tasks in two phases, while the master component 530A needs three phases to complete its task. Therefore, the slave components 530B, 530N are idle during an entire phase in which the corresponding master component 530A completes its respective task. The task coordinator 520 places the slave components 530B, 530N in a wait state during the phase in which the master component 530A completes its task (i.e. until the master component 530A sends back an "EndMethod" message to the task coordinator 520). Therefore, the operations of the master coordinator 530A and the slave coordinators 530B, 530N are synchronized with the "EndMethod" code.

In an alternative embodiment, the multi-phase protocol in accordance with the present invention can be implemented using components 530 that are not oriented in a master-slave format. In this embodiment, the components 530 control the phases and guide the task coordinator 520 to implement the multiple phase protocol using the multi-phase codes, such as NextPhase, SamePhase, and EndMethod. In this embodiment, the components 530 themselves can control whether they are to be invoked later by using the listed codes. A central control mechanism for parallel execution of tasks is thus not needed except for the task coordinator 520, which processes the codes to perform the requested invocation

procedures.

Figure 7 illustrates one embodiment of the method of performing parallel data operations in accordance with the present invention. The client system 120 receives a data transaction request from a user accessing the client system 120 (step 610). The client processor 430 receives through the user interface 440 and processes the data transaction request. The data transaction request may include data transactions such as data loading functions; data transformation functions, such as data selection, data validation, data cleansing and data condensing; data query functions; and the like. Once the client processor 430 receives the data transaction request, the client system 120 implements a multi-phase coordination protocol (step 620). A flow chart illustration of one embodiment of the steps in performing the multi-phase coordination protocol (step 620) is shown in Figure 8.

Figure 8 illustrates a flowchart depiction of one embodiment of the steps involved in preparing to perform the multi-phase coordination described above. The task coordinator 520 translates the task plan 595 (step 810 of Figure 8). Translating the plan comprises translating the functions indicated in the job script initiated by the client processor 430. Subsequently, the task coordinator 520 selects one or more components 530, or operators, in order to perform the operations as defined by the plan 595 (step 820). In one embodiment, the components 530 are software components plugged into the client infrastructure. The components 530 are capable of providing data extraction operations, data transformation operations, and data loading functions. In one embodiment, the client system 120 defines the first component 530A as the master process or the controlling instance. The client system 120 defines subsequent components 530 as the slave components, or the worker tasks.

The launching of the task coordinator 520 also comprises the step of generating a communications interface 590 (step 830). The client system 120

generates the communications interface 590 in such a manner that each component 530 can receive an input or send an output to other components 530 via the communications interface 590. The task coordinator 520 then coordinates the operation of the components 530 in a multi-phase parallel sequence (step 840). In one embodiment, the task coordinator 520 orchestrates the operation of the component 530 within a plurality of phases. In each phase, one or more components 530 may perform operations simultaneously, or in parallel, within the phase.

Figure 9 illustrates a flowchart representation of the steps involved in performing the multi-phase coordination protocol (step 620 in Figure 7) described above. The client controller 420 analyses the task plan 595 in order to determine whether a multi-phase processing is to be implemented (step 710). IN one embodiment, the task plan 595 comprises a job script generated by the client processor 430. When the client controller 420 makes a determination that multiphase processing is not to be implemented, the client system 120 enters a normal operation mode (step 720).

When the client controller 420 makes a determination that multiphase processing is to be implemented, the task coordinator 520 advances towards implementing the next component 530 by invoking the next method (step 730). In other words the task coordinator 520 invokes one or more components 530. The task coordinator 520 then increments a phase count by one (step 740). The task coordinator 520 invokes a component 530, by invoking a method, in the next phase, which corresponds to the incremented phase count (step 750).

The task coordinator 520 then analyzes the return code from the component 530 (step 760). When the component 530 returns a "SamePhase" message to the task coordinator 520, the task coordinator 520 again invokes the same component 530 within the same phase (see line 765). When the component 530 returns a "NextPhase" message to the task coordinator 520, the

task coordinator 520 increments the phase count by one and invokes the same component 530 in the next phase (see line 770). When the component 530 returns an "EndMethod" message to the task coordinator 520, the client controller 420 determines if multi-phase processing is to be implemented, and the steps shown in Figure 9 is repeated (see line 775). In an alternative embodiment, when the component 530 returns an "EndMethod" message to the task coordinator 520, the task coordinator 520 invokes the next component 530, and steps 730-760 are repeated.

Upon the completion of the steps described in Figure 9, the step of performing the multi-phase coordination protocol upon the data transaction requests (step 620), is substantially complete. In one embodiment, the client system 120 performs the multi-phase coordination protocol described in Figures 6, 7, 8, and 9, upon the database system 110. Turning back to Figure 7, once the transaction request is completed, the client system 120 is returned to normal operation (step 630).

Figure 10 illustrates one embodiment of the implementation of the client system 120, in accordance with the present invention. The system 900 is capable of implementing the multi-phase protocol described above. The system 900 comprises a network-attached UNIX client system 910. The network-attached UNIX client system 910 comprises a processor 945 and an operator interface 950. The processor 945 is capable of parsing, compiling, and executing a series of statements submitted through the job script 930. The job script 930 is based upon the data transaction request received from a user. In one embodiment, the INMOD routines 952 and the OUTMOD routines 954 are software routines that interface with a data source to allow for data load and data export functions, respectively.

The UNIX client system 910 comprises a plurality of operators 960A-960E. In one embodiment, the operators 960 are software components of the system

900 that provide data functions such as data extract, data transform, and data load functions. In one embodiment, the system 900 performs multi-phase parallel data operations upon data in a database, such as the RDBMS, which is a parallel relational database. The producer operators 960E are capable of data
5 extraction functions that include: getting data from the RDBMS 920; getting data from an external data source [not shown]; generate data internally; and pass data to other operators 960 via the operator interface 950. The consumer operators 960E are capable of data loading functions. The data loading functions performed by the consumer operator 960E include accepting data from other
10 operators 960 via the operator interface 950 and loading the data into the RDBMS 920, or loading the data to an external data source. The cleanse/filter/condensed operators 960D are capable of performing data transformation functions. The data transformation functions that can be performed by the cleanse/filter/condense operators 960D include data selection,
15 data validation, data cleansing, and data condensing.

In one embodiment, the operator interface 950 is a bi-directional call level interface that provides intercommunications between the processor 945 and the operators 960. The system 900 can load data and export data from any
20 accessible database object in the RDBMS 920 or other data source that contains an access operator and an access module. The destination of the data for functions for operations performed by the system 900 can be a relational database, a non-relational database, a database server, file, or a data storage device. The data that is processed by the system 900 may be text or objects
25 such as images, pictures, voice, or the like. Operations performed by the system 900 include retrieving, storing and transporting specific data objects via parallel data streams. The system 900 is also capable of performing merging or splitting multiple parallel data streams. Furthermore, the system 900 is capable of filtering, conditioning, and cleansing data.

30

The system 900 is capable of executing multiple instances of an operator. The system 900 has the ability to simultaneously read, transform, and load data. The parallel capabilities of the system 900 allows for data to be placed in a data buffer [not shown], while the next operator 960 can begin performing its tasks without waiting for the previous process to be completed. This capability allows for the elimination of intermediate storage devices since data is streamed through the processes performed by the system 900 without having the data written into storage disks.

The various devices and systems discussed each includes various software layers, routines, or modules. Such software layers, routines, or modules are executable on corresponding control units. The various control units in the devices or systems may each include a microprocessor, a microcontroller, a processor card (including one or more microprocessors or microcontrollers), or other control or computing devices. As used here, a "controller" refers to a hardware component, software component, or a combination of the two.

The storage devices referred to in this discussion may include one or more machine-readable storage media for storing data and instructions. The storage media may include different forms of memory including semiconductor memory devices such as dynamic or static random access memories (DRAMs or SRAMs), erasable and programmable read-only memories (EPROMs), electrically erasable and programmable read-only memories (EEPROMs) and flash memories; magnetic disks such as fixed, floppy and removable disks; other magnetic media including tape; and optical media such as compact disks (CDs) or digital video disks (DVDs). Instructions that make up the various software routines, modules, or layers in the various devices or systems may be stored in respective storage devices. The instructions when executed by a respective control unit cause the corresponding device or system to perform programmed acts.

The instructions of the software routines, modules, or layers may be loaded or transported to each device or system in one of many different ways. For example, code segments including instructions stored on floppy disks, CD or DVD media, a hard disk, or transported through a network interface card,

5 modem, or other interface device may be loaded into the device or system and executed as corresponding software routines, modules, or layers. In the loading or transport process, data signals that are embodied in carrier waves (transmitted over telephone lines, network lines, wireless links, cables, and the like) may communicate the code segments, including instructions, to the device
10 or system. Such carrier waves may be in the form of electrical, optical, acoustical, electromagnetic, or other types of signals.

While the invention has been disclosed with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and
15 variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of the invention.

What is claimed is:

1 1. A method of performing parallel data operations upon data in a
2 database, comprising:
3 receiving a data transaction request in a client system; and
4 executing a plurality of multi-phase parallel tasks in response to the
5 request to perform the data operations upon the data in the database.

1 2. The method of claim 1, wherein receiving a data transaction
2 request comprises receiving a request for loading data into the database.

1 3. The method of claim 1, wherein receiving a data transaction
2 request comprises receiving a request to perform a data transformation
3 operation upon the data in the database.

1 4. The method of claim 3, wherein receiving a request to perform the
2 data transformation operation comprises receiving a request to perform one of a
3 data selection operation, a data validation operation, a data cleansing operation,
4 and a data query operation.

1 5. The method of claim 1, wherein executing the multi-phase parallel
2 tasks comprises executing each of the parallel tasks in one or more phases.

1 6. The method of claim 5, comprising executing a first parallel task in
2 a first number of phases and a second parallel task in a second number of
3 phases.

1 7. The method of claim 5, further comprising each parallel task
2 providing a code to indicate if the task is to be re-invoked in the next phase

8 creating a communications channel to allow for communications
9 between the processes.

1 14. The method of claim 13, wherein selecting the plurality of
2 components to implement the translated task plan comprises selecting the
3 plurality of components to perform at least one of a data extraction operation, a
4 data transformation operation, and a data loading operation.

1 15. An apparatus, comprising:
2 a user interface;
3 a processor coupled with the user interface, wherein the processor
4 receives a data transaction request from the user interface; and
5 a controller coupled with the processor, wherein the controller
6 performs a number of tasks in parallel based upon instructions received from the
7 processor, each task performed in a plurality of phases.

1 16. The apparatus of claim 15, wherein the processor generates a task
2 plan in response to the data transaction request.

1 17. The apparatus of claim 16, wherein the controller comprises a task
2 coordinator to execute the task plan.

1 18. The apparatus of claim 16, wherein the controller further comprises
2 a plurality of components to implement the task plan in parallel.

1 19. The apparatus of claim 18, further comprising a communications
2 interface for enabling communications between the components.

1 20. The apparatus of claim 18, wherein the controller further comprises
2 a storage unit for storing methods and functions to execute the task plan.

1 22. A system, comprising:
2 a database system; and
3 a client system to implement a plurality of data operations upon the
4 database in parallel.

1 23. The system of claim 22, wherein the database system is a relational
2 database system.

1 24. The system of claim 23, wherein the database system is a parallel
2 database system.

1 25. The system of claim 22, wherein client system comprises:
2 a processor to receive a data transaction request;
3 a plurality of operators to perform parallel data operations in
4 response to the data transaction request;
5 an operator interface coupled to the operators, wherein the
6 operator interface allows communications between the operators.

1 26. The system of claim 22, wherein the processor performs data
2 parsing and data compiling upon the data in the database system.

1 27. The system of claim 22, wherein the operators perform at least one
2 of a data extraction function, a data transform function, and a data loading
3 function.

1 28. An article comprising at least one storage medium containing
2 instructions that when executed cause a client system to:
3 receive a data transaction request; and
4 execute a plurality of parallel tasks in response to the request to
5 perform data operations upon the data in the database over a network
6 connection.

1 29. The article of claim 28, wherein the instructions when executed
2 cause the client system to execute each of the parallel tasks in one or more
3 phases.

1 30. The article of claim 29, wherein the instructions when executed
2 cause the client system to execute a first parallel task in a first number of phases
3 and a second parallel task in a second number of phases.

1 31. The article of claim 29, wherein the instructions when executed
2 cause each parallel task to provide a code to indicate if the task is to be re-
3 invoked in the next phase.

1 32. The article of claim 31, wherein the instructions when executed
2 cause the parallel task to provide the code to a task coordinator.

1 33. The article of claim 32, wherein the code comprises a first code to
2 indicate that the task coordinator is to invoke a component in the next phase.

1 34. The article of claim 32, wherein the code comprises a second code
2 to indicate that the task is not to invoke the component in the next phase.

1 35. A method of performing parallel data operations upon data in a
2 database, comprising:
3 receiving a data transaction request; and

4 executing a plurality of synchronized multi-phase parallel tasks in
5 response to the request to perform the data operations upon the data in the
6 database.

1 36. The method of claim 35, wherein executing the multi-phase parallel
2 tasks comprises executing each of the parallel tasks in one or more phases.

1 37. The method of claim 36, comprising executing a first parallel task in
2 a first number of phases and a second parallel task in a second number of
3 phases.

1 38. The method of claim 36, further comprising each parallel task
2 providing a code to indicate if the task is to be re-invoked in the next phase.

1 39. The method of claim 38, wherein providing the code comprises
2 providing the code to a task coordinator.

1 40. The method of claim 39, wherein the code comprises a first code to
2 indicate that the task coordinator is to invoke a component in the next phase.

1 41. The method of claim 39, wherein the code comprises a second
2 code to indicate that the task is not to invoke a component in the next phase.

1 42. The method of claim 39, wherein the code synchronizes the
2 operation of one or more component.

ABSTRACT OF THE INVENTION

A method and apparatus enables a client system to perform multi-phase parallel operations upon data in a database system. A data transaction request received by the client system. A plurality of multi-phase parallel tasks are executed in response to the request to perform the data operations upon the data in the database.

006650" 2983/350




FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 4




FIGURE 5

Methods	Phase	Master	Slave (1)	Slave (2)
Startup	1	Allocate Global Var	Create/Connect an SQL session	Create/Connect an SQL session
		Create/Connect an SQL session.	NextPhase	NextPhase
		NextPhase		
	2	Create/Connect Fastload session(s).	Create/Connect Fastload session(s).	Create/Connect Fastload session(s).
		NextPhase	EndMethod	EndMethod
Execute	3	Issue 'Begin Loading'		
		EndMethod		
	1	Read a record from a data source	Read a record from a data source	Read a record from a data source
		IF End-Of-Data NextPhase ELSE SamePhase	IF End-Of-Data NextPhase ELSE SamePhase	IF End-Of-Data NextPhase ELSE SamePhase
	2	Issue 'End Loading' (if specified by the user)		
Check-point		EndMethod		
	1	Write statistics to Global Var	Write statistics to Global Var	Write statistics to Global Var
		NextPhase	EndMethod	EndMethod
	2	Write checkpoint record(s) to the LOG		
		EndMethod		

FIGURE 6A

Methods	Phase	Master	Slave (1)	Slave (2)
Restart	1	Process checkpoint records	Process checkpoint records	Process checkpoint records
		NextPhase	EndMethod	EndMethod
	2	Write a restart record to the LOG to indicate EndOfRestart		
		EndMethod		
Terminate	1	Disconnect Fastload session(s)	Disconnect Fastload session(s)	Disconnect Fastload session(s)
		NextPhase	NextPhase	NextPhase
	2	Disconnect SQL session	Disconnect SQL session	Disconnect SQL session
		Deallocate GlobalVar	EndMethod	EndMethod
		EndMethod		

FIGURE 6B

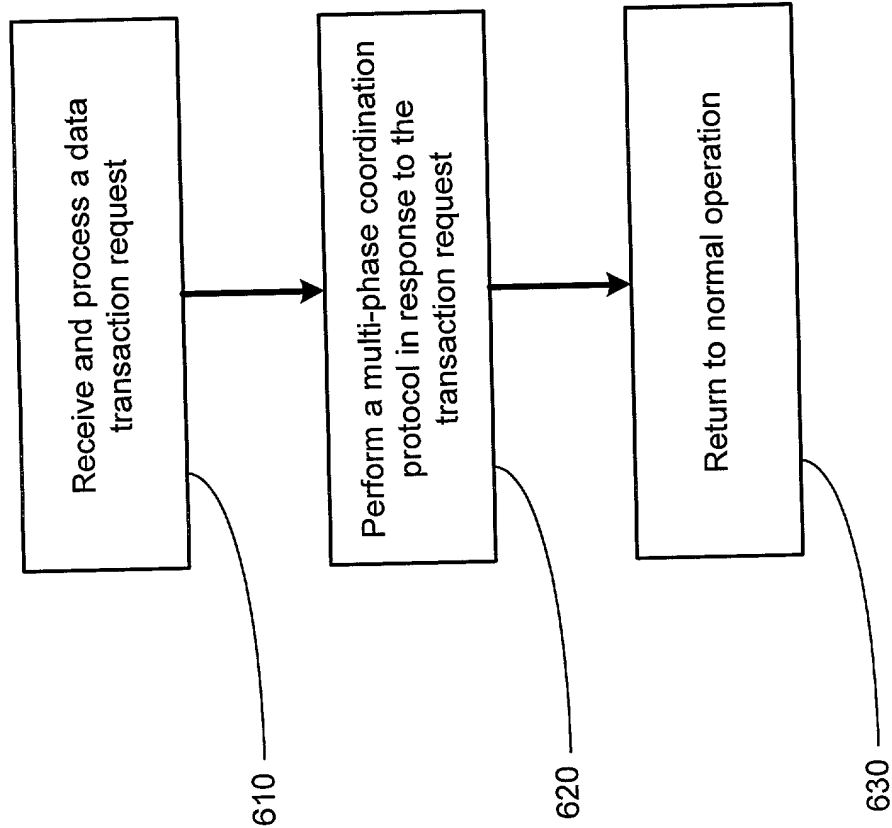


FIGURE 7

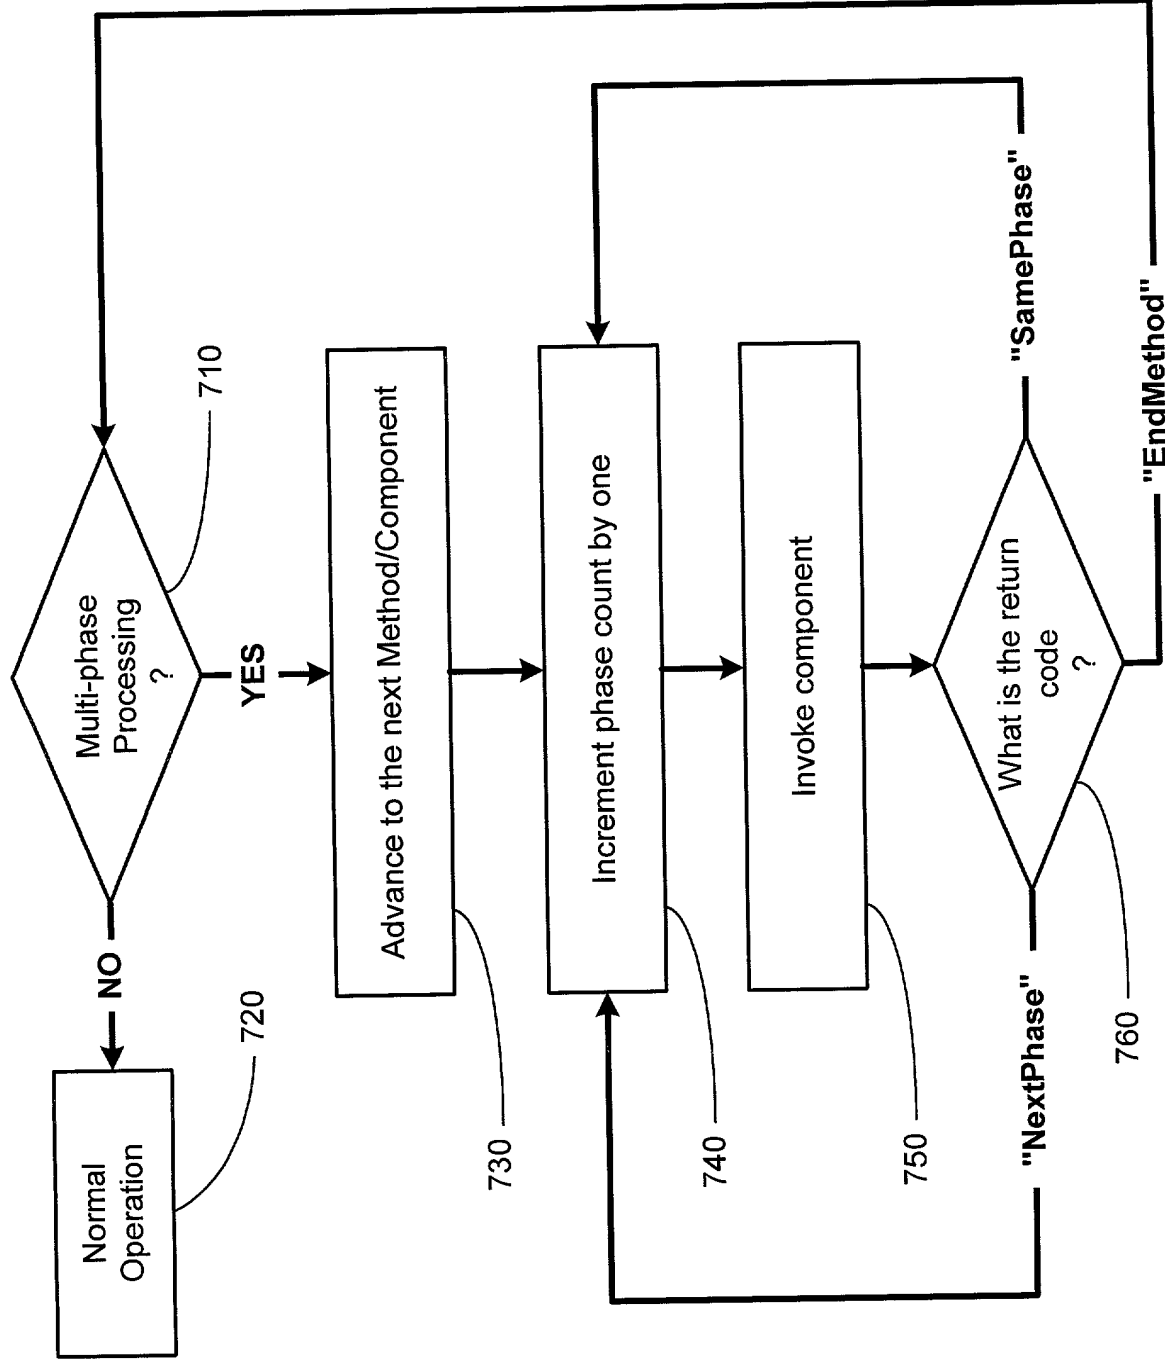


FIGURE 9





FIGURE 10

United States Patent & Trademark Office
Office of Initial Patent Examination -- Scanning Division



SCANNED. # 6

Application deficiencies were found during scanning:

☒ Page(s) _____ of Express Mailing were not present
for scanning. (Document title)

☐ Page(s) _____ of _____ were not present
for scanning. (Document title)

☐ Scanned copy is best available.
